Single Channel Audio Source Separation using Convolutional Denoising Autoencoders

Emad M. Grais and Mark D. Plumbley

Centre for Vision, Speech and Signal Processing, University of Surrey, Guildford, UK.

{grais, m.plumbley}@surrey.ac.uk

arXiv:1703.08019v1 [cs.SD] 23 Mar 2017

Abstract

Deep learning techniques have been used recently to tackle the audio source separation problem. In this work, we propose to use deep convolution denoising auto-encoders (CDAEs) for monaural audio source separation. We use as many CDAEs as the number of sources to be separated from the mixed signal. Each CDAE is trained to separate one source and treats the other sources as background noise. The main idea is to allow each CDAE to learn suitable time-frequency filters and features to its corresponding source. Our experimental results show that CDAEs perform source separation slightly better than the deep feedforward neural networks (FNNs) even with a much less number of parameters than FNNs.

Index Terms: Convolutional denoising autoencoders, single channel audio source separation, stacked convolutional autoencoders, deep convolutional neural networks, deep learning.

1. Introduction

Different types of deep neural networks (DNNs) have been used to tackle the audio single channel source separation (SCSS) problem [1, 2, 3, 4]. The denoising auto-encoder (DAE) is a special type of fully connected feedforward neural networks that takes noisy input signals and outputs their denoised version [5, 6]. DAEs are common in deep learning, they are used to learn robust low-dimensional features even when the inputs are perturbed with some noise [7, 8]. DAEs have been used for SCSS where the inputs of the DAE are the spectral frames of the mixed signal and the outputs are the spectral frames of the target source [3, 9]. Fully connected DAEs can not capture the 2D (spectral-temporal) structures of the spectrogram of the input and output signals. Since DAEs are fully connected networks, they usually have a lot of parameters to be optimized.

For their ability in extracting robust spectral-temporal structures of different audio signals [10], convolutional neural networks (CNN) have been used successfully to learn useful features in many audio processing applications such as: speech recognition [11], speech enhancement [12], audio tagging [13], and many music related applications [14, 15, 16]. Convolutional denoising auto-encoders (CDAEs) are also a special type of CNNs that can be used to discover robust localized low-dimensional patterns that repeat themselves over the input [17, 18]. CDAEs differ from conventional DAEs as their parameters (weights) are shared which makes the CDAEs have much less parameters than DAEs. The ability of CDAEs to extract repeating patterns in the input makes them suitable to be used to suppress the background music signal that is embedded with speech signals to improve the quality of the speech recognition system [19]. The results in [19] show that CDAE works better than DAE in removing the background music signals from the speech signal. Recently, a fully convolutional denoising auto-encoders (all the network layers are composed of convolutional units) were used for speech enhancement [20] that maps the distorted speech signal to its clean speech signal. The results in [20] show that CDAE works as well as the fully connected feedforward neural networks (FNNs) and recurrent neural networks (RNNs) even when the number of parameters in the CDAE is much less than the number of parameters in FNNs and RNNs.

Motivated by the aforementioned successes of using CNNs and CDAEs in a variety of audio signals, we propose in this paper to use deep (stacked) fully convolutional DAEs (CDAEs) for the audio single channel source separation (SCSS). The main idea in this paper is to train a CDAE to extract one target source from the mixture and treats the other sources as background noise that needs to be suppressed. This means we need as many CDAEs as the number of sources that need to be separated from the mixed signal. This is a very challenging task because each CDAE has to deal with highly nonstationary background signals/noise. Each CDAE sees the magnitude spectrograms as 2D segments which helps in learning the temporal and spectral information for the audio signals. From the ability of CDAEs in learning noise robust features, in this work, we train each CDAE to learn unique temporal-spectral patterns for its corresponding target source. Each trained CDAE is then used to extract/separate the related patterns of its corresponding target source from the mixed signal.

This paper is organized as follows: Section 2 shows a brief introduction about CDAEs. The proposed approach of using CDAEs for SCSS is presented in Section 3. The experiments and discussions are shown in Section 4.

2. Fully convolutional denoising autoencoders

Convolutional auto-encoders (CAEs) [17, 18] are a special type of CNNs. As shown in Fig. 1, CAE is composed of two main parts, the encoder part and the decoder part. The encoder part maps the input data into low dimensional features. The decoder part reconstructs the input data from the low dimensional features.

Convolutional denoising auto-encoders (CDAEs) are similar to CAEs but CDAEs are trained from corrupted input signals and the encoder is used to extract noise robust features that the decoder can use to reconstruct a cleaned-up version of the input data [19, 20]. The encoder part in CDAEs is composed of repetitions of a convolutional layer, an activation layer, and a pooling layer as shown in Fig. 1. The convolutional layers consist of a set of filters that extract features from their input layers, the activation layer in this work is the rectified linear unit that imposes nonlinearity to the feature maps. The pooling in this work is chosen to be max-pooling [21]. Maxpooling does the downsampling of the latent representation by a constant factor taking the maximum value within a certain scope of the mapping space and generates new mapping space with a reduced dimension.

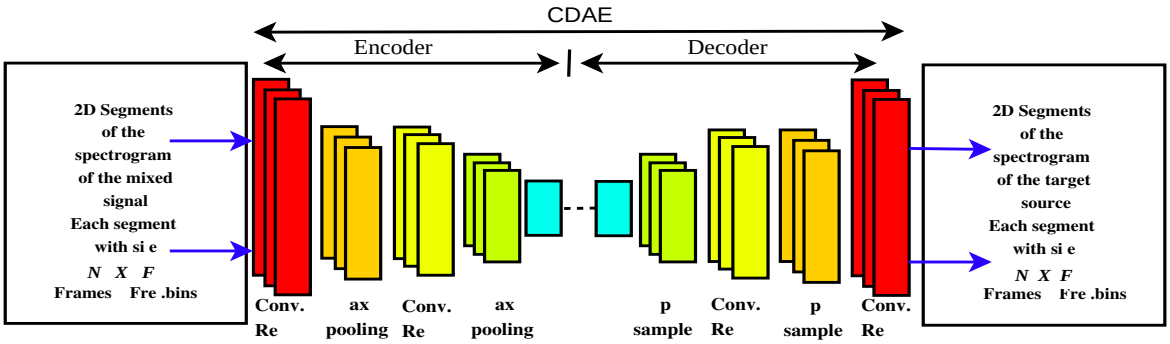

Figure 1: *The overview of the proposed structure of a convolutional denoising auto-encoder (CDAE) that separates one target source from the mixed signal. "Conv." denotes a 2D convolutional layer, ReLU denotes a rectified linear unit as an activation function. We use a CDAE for each source.*

The final goal of the encoder part is to extract noise robust low dimensional features from the input data. The max-pooling is the layer that reduces the dimensionality of the mapping space. This can be seen as 2D nonlinear low rank representations for the input data (considering that the data in this paper are 2-D signals). The decoder part consists of repetitions of a convolutional layer, an activation layer, and an upsampling layer. The up-sampling layer does the up-sampling on the feature maps of the previous layer and generates new ones with high dimension. More details about CAE can be found in [17, 18]. In this work, the data are 2D signals (magnitude spectrograms), filters, pooling, and upsampling are all 2D operators.

3. The proposed approach of using CDAEs for SCSS

Given a mixture of I sources as $y(t) = \sum_{i=1}^{I} s_i(t)$, the aim of audio SCSS is to estimate the sources $s_i(t)$, $\forall i$, from the mixed signal $y(t)$. We work here in the short time Fourier transform (STFT) domain. Given the STFT of the mixed signal $y(t)$, the main goal is to estimate the STFT of each source in the mixture.

The ability of the CDAE in denoising motivates us to use it for source separation where each source is considered as a target signal and all the other sources in the mixture are considered to be background noise. This is a challenging task for each CDAE since each CDAE deals with highly nonstationary background noise (other sources in the mixture). In this paper, we propose to use as many CDAEs as the number of sources to be separated from the mixed signal. Each CDAE sees the mixed signal as a combination of a target source and background noise. The main aim of each CDAE is to estimate a clean signal for its corresponding source from the other background sources that exist in the mixed signal. Each CDAE is trained to map the magnitude spectrogram of the mixture into the magnitude spectrogram of one of the target sources. Each CDAE in this work is a fully 2D convolutional deep neural network without any fully connected or recurrent layers, which keeps the number of parameters to be optimized for each CDAE very small. Each CDAE is composed of encoder and decoder networks. The encoder part aims to extract noise robust low denominational features and the decoder part aims to reconstruct a clean signal from the low dimensional features [19, 20]. The inputs and outputs of the CDAEs are 2D-segments from the magnitude spectrograms of the mixed and target signals respectively. Therefore, the CDAEs span multiple time frames to capture the time-frequency characteristics of each source. The number of frames that each input segment has is N and the number of frequency bins is F. In this work, F is the dimension of the whole spectral frame.

3.1. Training the CDAEs for source separation

Let's assume we have training data for the mixed signals and their corresponding clean sources. Let $\mathbf{Y}_{\text{tr}}$ be the magnitude spectrogram of the mixed signal and $\mathbf{S}_i$ be the magnitude spectrogram of the clean source i. The subscript "tr" denotes the training data. The CDAE of source i is trained to minimize the following cost function:

$$C_i = \sum_{n,f} \left(\mathbf{Z}_i\left(n,f\right) - \mathbf{S}_i\left(n,f\right)\right)^2 \tag{1}$$

where $\mathbf{Z}_i$ is the actual output of the last layer of the CDAE of source i, $\mathbf{S}_i$ is the reference clean output signal for source i, n, and f are the time and frequency indices respectively. The input of all the CDAEs is the magnitude spectrogram $\mathbf{Y}_{\text{tr}}$ of the mixed signal.

The input and output instants of the CDAEs are 2D-segments from the spectrograms of the mixed and target signals respectively. Each segment is composed of N consequent spectral frames taken from the magnitude spectrograms. This allows the CDAEs to learn unique time-frequency patterns for each source.

3.2. Testing the CDAEs for source separation

Given the trained CDAEs, the magnitude spectrogram $\mathbf{Y}$ of the mixed signal is passed through all the trained CDAEs. The output of the CDAE of source i is the estimate of the spectrogram $\tilde{\mathbf{S}}_\mathbf{i}$ of source i.

4. Experiments and Discussion

We applied our proposed single channel source separation (SCSS) using CDAEs approach to separate audio sources from a group of songs from the SiSEC-2015-MUS-task dataset [22]. The dataset has 100 stereo songs with different genres and instrumentations. To use the data for the proposed SCSS approach, we converted the stereo songs into mono by computing the average of the two channels for all songs and sources in the data set. Each song is a mixture of vocals, bass, drums, and other musical instruments. We used our proposed algorithm to separate each song into vocals, bass, drums and other instruments. The other instruments (*other* for short) here (sources that are not vocal, bass, or drums) are treated as one source. We trained four CDAEs for the four sources (vocals, bass, drums, and other).

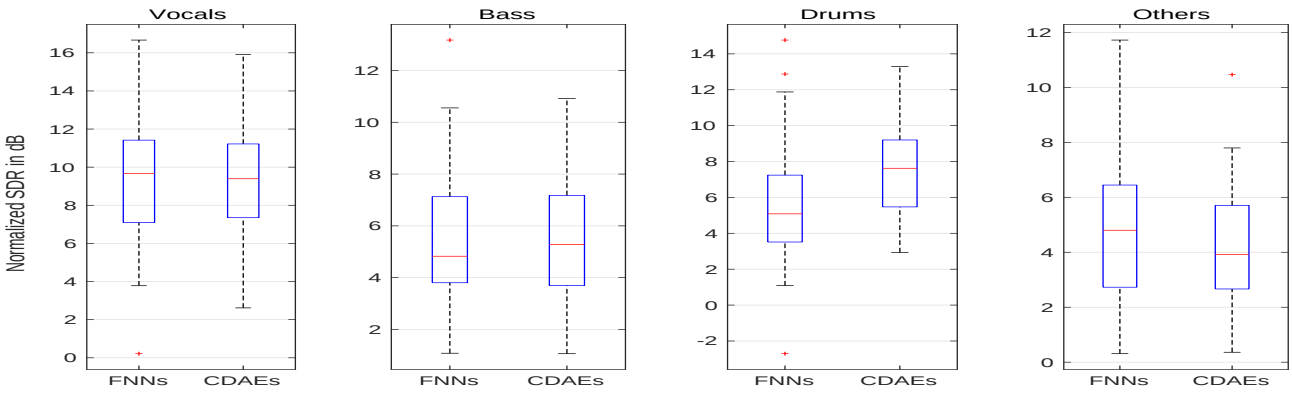

Figure 2: *The normalized SDR values in dB of using fully connected feedforward neural networks (FNNs) and the proposed convolutional denoising autoencoders (CDAEs) method for source separation.*

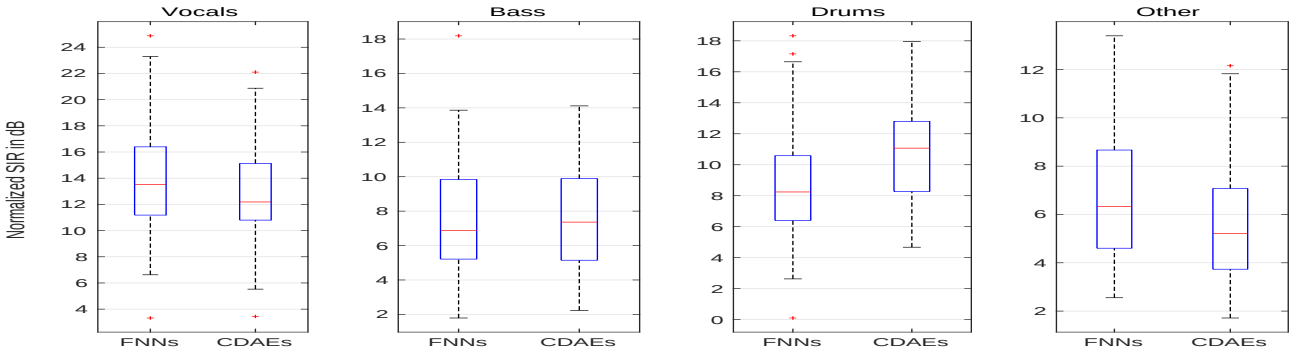

Figure 3: *The normalized SIR values in dB of using fully connected feedforward neural networks (FNNs) and the proposed convolutional denoising autoencoders (CDAEs) method for source separation.*

The first 50 songs were used to train all the networks for separation, and the last 50 songs were used for testing. The data was sampled at 44.1kHz sampling rate. The magnitude spectrograms for the data were calculated using the STFT, a Hanning window with 2048 points length and overlap interval of 512 was used and the FFT was taken at 2048 points, the first 1025 FFT points only were used since the conjugate of the remaining 1024 points are involved in the first points.

For the input and output data for the CDAEs, we chose the number of spectral frames in each 2D-segment to be 15 frames. This means the dimension of each input and output instant for each CDAE is 15 (time frames) $\times$ 1025 (frequency bins). Thus, each input and output instants (the 2D-segments from the spectrograms) spans around 370 msec of the waveforms of the data.

For the CDAES, each CDAE has seven hidden 2D convolutional layers with rectified linear unit (Relu) as an activation function. The details of the dimensions of the input, convolutional, max-pooling, up-sampling and output layers of each CDAE are shown in Table 1. The dimensions are shown as time-frames $\times$ frequency. The size of each filter is 3$\times$ 3 as in [11, 16]. "MaxPooling2D(3,5)" in Table 1, denotes downsampling the feature maps by 3 in the time-frame direction and by 5 in the frequency direction of the 2D feature maps. "UpSampling2D(3,5)" in Table 1, denotes upsampling the feature maps by 3 in the time-frame direction and by 5 in the frequency direction. The output layer is also a convolutional layer with Relu that its size was adjusted to match the size of each 2D output segment (15$\times$1025) of the spectrogram. As can be seen from Table 1, each CDAE has 37101 parameters.

We compared our proposed SCSS using CDAEs approach for SCSS with using deep fully connected feedforward neural networks (FNN) for SCSS approach. Four FNNs were used and each FNN was used to separate one source. Each FNN has three hidden layers with rectified linear units as activation functions. Each hidden layer has 1025 nodes. The parameters of the networks are tuned based on our previous work on the same dataset [23, 24]. The number of parameters in each FNN is 4206600 parameters which is greater than 112 times the number of parameters in each CDAE.

The parameters for all the networks were initialized randomly. All the networks were trained using backpropagation with Nesterovs accelerated gradient [25] as recommended by [26] with parameters: $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-08$, schedule-decay$= 0.004$, batch size 100, and a learning rate starts with 0.002 and reduced by a factor of 10 when the values of the cost function do not decrease on the validation set for 3 consecutive epochs. The maximum number of epochs was 100. We implemented our proposed algorithm using Keras based on Theano [27, 28].

Since we separate all four sources from the mixed signals, it is preferred to separate the sources from the mixed signal by building spectral masks that scale the mixed signal according to the contribution of each source in the mixed signal [29, 30]. The masks make sure that the sum of the estimated sources adds up to the mixed signal. Here we used the output spectrograms $\tilde{\mathbf{S}}_\mathbf{i}$, $\forall i$ of the networks to build spectral masks as follows:

$$\mathbf{M}_i(n,f) = \frac{\tilde{\mathbf{S}}_\mathbf{i}(n,f)}{\sum_j^I \tilde{\mathbf{S}}_\mathbf{j}(n,f)}, \ \forall i. \tag{2}$$

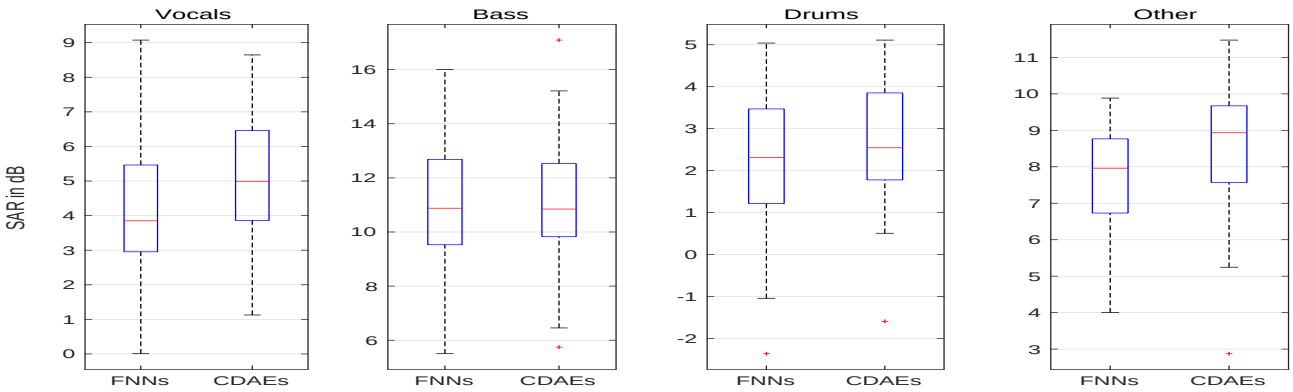

Figure 4: *The SAR values in dB of using fully connected feedforward neural networks (FNNs) and the proposed convolutional denoising autoencoders (CDAEs) method for source separation.*

The final estimate for source i can be found as

$$\hat{\mathbf{S}}_\mathbf{i}(n, f) = \mathbf{M}_i(n, f) \times \mathbf{Y}(n, f) \tag{3}$$

where $\mathbf{Y}(n, f)$ is the magnitude spectrogram of the mixed signal at frame n and frequncy f. The time domain estimate for source $\hat{s}_i(t)$ is computed using the inverse STFT of $\hat{\boldsymbol{S}}_i$ with the phase angle of the STFT of the mixed signal.

CDAE model summary		
The input data with size 15 frames and 1025 frequency bins		
Layer (type)	Number of filters	output shape
Convolution2D(3,3) MaxPooling2D(3,5)	12	(15, 1025) (5, 205)
Convolution2D MaxPooling2D(1,5)	20	(5, 205) (5, 41)
Convolution2D	30	(5, 41)
Convolution2D	40	(5, 41)
Convolution2D	30	(5, 41)
Convolution2D UpSampling2D(1,5)	20	(5, 41) (5, 205)
Convolution2D(3,3) UpSampling2D(3,5)	12	(5, 205) (15, 1025)
Convolution2D	1	(15, 1025)
The output data with size 15 frames and 1025 frequency bins Total number of parameters: 37101		

Table 1: *The detail structure of each CDAE. The output shape is shown as (time-frame , frequency). "Convolution2D(3,3)" denotes 2D convolutional layer with filter size 3×3. "MaxPooling2D(3,5)" denotes downsampling by 3 in the time-frame direction and by 5 in the frequency direction. "UpSampling2D(3,5)" denotes upsampling by 3 in the time-frame direction and by 5 in the frequency direction.*

Performance of the separation algorithms was measured using the signal to distortion ratio (SDR), signal to interference ratio (SIR), and signal to artefact ratio (SAR) [31]. SIR indicates how well the sources are separated based on the remaining interference between the sources after separation. SAR indicates the artefacts caused by the separation algorithm to the estimated separated sources. SDR measures how distorted the separated sources are. The SDR values are usually considered as the overall performance evaluation for any source separation approach [31]. Achieving high SDR, SIR, and SAR indicates good separation performance.

Figs. 2 and 3 show the box-plots of the normalized SDR and SIR values respectively. The normalization was done by subtracting the SDR and SIR values of the mixed signal from their corresponding values of the estimated sources [32]. Fig. 4 shows the box-plots of the SAR values of the separated sources using FNNs and CDAEs. We show the SAR values without normalization because SAR for the mixed signal is almost infinity.

As can be seen from Figs. 2, 3, and 4, CDAEs work significantly better than FNNs on the drums, which is considered as a difficult source to be separated [1, 29]. For the SDR and SIR, the performance of FNNs and CDAEs is almost the same for the vocals and bass sources. FNNs performs better than CDAEs on the other sources' case. The other sources here represent all sources that appear in the songs that are not vocals, drums, and bass, which is a variety of sources that change from one song to another. For SAR, the performance of CDAEs and FNNs is the same for bass, but CDAEs perform better than FNNs in the remaining sources. In general, it is not easy to have a fair comparison between the two methods since FNNs have more parameters than CDAEs and CDAEs were applied on time-frequency segments (without overlapping) while the FNNs were applied on individual spectral frames.

5. Conclusions

In this work we proposed a new approach for single channel source separation (SCSS). The new approach is based on using deep (stacked) fully convolutional denoising autoencoders (CDAEs). We used as many CDAEs as the number of sources to be separated from the mixed signal. Each CDAE learns unique patterns for each source and uses this information to separate the related components of each source from the mixed signal. The experimental results indicate that using CDAEs for SCSS is a promising approach and with very few parameters can achieve competitive results with the feedforward neural networks.

In our future work, we will investigate the effect of the following CDAE's parameters on the quality of the separated sources: the size of the filters, the number of the filters in every layer, the maxpooling and upsampling ratios, the number of frames in every input/output segment, and the number of layers. We will also investigate the possibility of using CDAEs for the multi-channel audio source separation problem using multi-channel convolutional layers.

6. Acknowledgements

This work is supported by grants EP/L027119/1 and EP/L027119/2 from the UK Engineering and Physical Sciences Research Council (EPSRC).

7. References

[1] S. Uhlich, M. Porcu, F. Giron, M. Enenkl, T. Kemp, N. Takahashi, and Y. Mitsufuji, "Improving music source separation based on deep neural networks through data augmentation and network blending," in *Proc. ICASSP*, 2017.

[2] E. M. Grais, G. Roma, A. J. Simpson, and M. D. Plumbley, "Discriminative enhancement for single channel audio source separation using deep neural networks," in *Proc. LVA/ICA*, 2017, pp. 236–246.

[3] M. Kim and P. Smaragdis, "Adaptive denoising autoencoders: A fine-tuning scheme to learn from test mixtures," in *Proc. LVA/ICA*, 2015, pp. 100–107.

[4] P. Chandna, M. Miron, J. Janer, and E. Gomez, "Monoaural audio source separation using deep convolutional neural networks," in *Proc. LVA/ICA*, 2017, pp. 258–266.

[5] J. Xie, L. Xu, and E. Chen, "Image denoising and inpainting with deep neural networks," in *Advances in NIPS*, 2012.

[6] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P. A. Manzagol, "Stacked Denoising Autoencoders: learning useful representations in a deep network with a local denoising criterion," *Journal of Machine Learning Research*, vol. 11, pp. 3371 – 3408, 2010.

[7] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proc. ICML*, 2008, pp. 1096–1103.

[8] G. Hinton and R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504 – 507, 2006.

[9] P. Smaragdis and S. Venkataramani, "A neural network alternative to non-negative audio models," in *Proc. ICASSP*, 2017.

[10] H. Lee, P. Pham, Y. Largman, and A. Y. Ng, "Unsupervised feature learning for audio classification using convolutional deep belief networks," in *Advances in NIPS*, 2009, pp. 1096–1104.

[11] Y. Qian, M. Bi, T. Tan, and K. Yu, "Very deep convolutional neural networks for noise robust speech recognition," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 24, no. 12, pp. 2263–2276, 2016.

[12] S.-W. Fu, Y. Tsao, and X. Lu, "SNR-Aware convolutional neural network modeling for speech enhancement," in *Proc. InterSpeech*, 2016.

[13] Y. Xu, Q. Kong, Q. Huang, W. Wang, and M. D. Plumbley, "Convolutional gated recurrent neural network incorporating spatial features for audio tagging," in *Proc. IJCNN*, 2017.

[14] Y. Han, J. Kim, and K. Lee, "Deep convolutional neural networks for predominant instrument recognition in polyphonic music," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 208–221, 2017.

[15] K. Choi, G. Fazekas, M. Sandler, and K. Cho, "Convolutional recurrent neural networks for music classification," in *arXiv:1609.04243*, 2016.

[16] F. Korzeniowski and G. Widmer, "A fully convolutional deep auditory model for musical chord recognition," in *Proc. Workshop on Machine Learning for Signal Processing*, 2016.

[17] J. Masci, U. Meier, D. Ciresan, and J. Schmidhuber, "Stacked convolutional auto-encoders for hierarchical feature extraction," in *Advances in NIPS*, 2011.

[18] B. Du, W. Xiong, J. Wu, L. Zhang, L. Zhang, and D. Tao, "Stacked convolutional denoising auto-encoders for feature representation," *IEEE Trans. on Cybernetics*, vol. 47, no. 4, pp. 1017–1027, 2017.

[19] M. Zhao, D. Wang, Z. Zhang, and X. Zhang, "Music removal by convolutional denoising autoencoder in speech recognition," in *In proc. APSIPA*, 2016.

[20] S. R. Park and J. W. Lee, "A fully convolutional neural network for speech enhancement," in *arXiv preprint arXiv:1609.07132*, 2016.

[21] D. Scherer, A. Muller, and S. Behnke, "Evaluation of pooling operations in convolutional architectures for object recognition," in *Advances in NIPS*, 2010.

[22] N. Ono, Z. Rafii, D. Kitamura, N. Ito, and A. Liutkus, "The 2015 signal separation evaluation campaign," in *Proc. LVA/ICA*, 2015, pp. 387–395.

[23] E. M. Grais, G. Roma, A. J. R. Simpson, and M. D. Plumbley, "Single channel audio source separation using deep neural network ensembles," in *Proc. 140th Audio Engineering Society Convention*, 2016.

[24] ——, "Combining mask estimates for single channel audio source separation using deep neural networks," in *Prec. InterSpeech*, 2016.

[25] Y. Nesterov, "A method of solving a convex programming problem with convergence rate o(1/sqr(k))," *Soviet Mathematics Doklady*, vol. 27, no. 2, pp. 372–376, 1983.

[26] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, "On the importance of initialization and momentum in deep learning," in *Proc. ICML*, vol. 28, no. 3, May 2013, pp. 1139–1147.

[27] F. Chollet, "Keras," https://github.com/fchollet/keras, 2015.

[28] F. Bastien, P. Lamblin, R. Pascanu, J. Bergstra, I. Goodfellow, A. Bergeron, N. Bouchard, D. Warde-Farley, and Y. Bengio, "Theano: new features and speed improvements," in *Deep learning and unsupervised feature learning NIPS workshop*, 2012.

[29] A. A. Nugraha, A. Liutkus, and E. Vincent, "Multichannel audio source separation with deep neural networks," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 24, no. 9, pp. 1652–1664, 2016.

[30] P.-S. Huang, M. Kim, M. Hasegawa-Johnson, and P. Smaragdis, "Singing-Voice separation from monaural recordings using deep recurrent neural networks," in *Proc. ISMIR*, 2014, pp. 477–482.

[31] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 14, no. 4, pp. 1462–69, Jul. 2006.

[32] A. Ozerov, P. Philippe, F. Bimbot, and R. Gribonval, "Adaptation of Bayesian models for single-channel source separation and its application to voice/music separation in popular songs," *IEEE Trans. of Audio, Speech, and Language Processing*, vol. 15, pp. 1564–1578, 2007.